

CONTEXT-SENSITIVE CONSTRAINT DRIVEN UNIQUIFICATION AND CHARACTERIZATION OF STANDARD CELLS

5 BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to a method and system for optimizing the creation, utilization, and characterization of standard cells used in the design of integrated circuits, that are particularly suitable for design-specific standard cells.

2. Description of the Related Art

Integrated circuits (ICs) form the heart of electronic engineering technology. The design and fabrication processes of ICs are extremely complex. Numerous methodologies, based on various technologies and tools, are known for accomplishing the design and fabrication of ICs. Prominent IC design methodologies include custom, semi-custom, gate-array, standard-cell, FPGA (field-programmable gate arrays).

It is generally accepted that standard cell based solutions provide an acceptable balance in terms of cost and performance for the implementation of ASICs (application specific integrated circuits), as well as several other forms of ICs. A typical standard-cell based design is implemented using a specific set of predefined logic cells referred to as standard-cells. The set of standard-cells is typically referred to as a library. Each (unique) standard cell is typically instantiated multiple times and the IC design is realized by the creation of appropriate interconnections between the instantiations.

Popular methods for implementing a standard cell based IC design use a fixed library of standard-cells for multiple designs. While a typical state-of-the-art standard-cell library may contain as many as several hundred to a thousand or more cells, IC design synthesis tools tend to use a small subset of the available cells – typically in the range of 80-200 unique cells.

Consequently, there is a great need to efficiently identify the cells in a given library that may be able to provide a better implementation of the functionality of a design, as measured using commonly used metrics such as performance or area.

5 More recently, digital IC design methodologies have been developed that utilize on-the-fly generation of standard cells as disclosed in commonly assigned U.S. patent serial number 09/896071, entitled Process for Automated Generation of Design-Specific Complex Functional Blocks to Improve Quality of Synthesized Digital Integrated Circuits in CMOS, and incorporated herein by reference. In order to improve the efficiency of such design
10 methodologies, there exists the need to (1) reduce the number of unique cells required to implement the design (Uniquification); and (2) efficiently characterize the properties of newly implemented cells (Characterization).

A primary focus in conventional, automated, design methodologies is placed on
15 effective technology mapping. That is, the process of mapping the design to an existing technology-dependent library of standard-cells. The most important step in the mapping process is that of matching a target functional description with the functionality of a physical implementation (e.g., a standard-cell). For example, consider the example of Fig. 1. The target function $f = (xy + zy)'$ can be implemented by available AOI cell 105 in the library
20 whose functionality is given by $f = (a(b + c))'$. However, in order to achieve this implementation, the inputs of the target function would need to be permuted in a certain order. In this example, two such permutations can result in an implementation of the target function with AOI cell 105. The two permutations that can result in a match in functionality are (a:y, b:x, c:z) and (a:y, b:z, c:x), as shown in the table of Fig. 1a.

25

It should also be noted that a function $f = (xy + z'y)'$ can also be realized by the same AOI cell 105, if inverters were permitted to be used freely in the mapping process. Both of the permutations introduced above would still be valid, except that an inverter is required at the input of the AOI cell that matches the z input of the target cell in order to realize the

correct target function. For example, in the permutation (a:y, b:x, c:z), the input c requires an inverter in order to realize the target function.

The technology mapping problem is well-known in conventional logic synthesis, it belongs to the category of computationally difficult problems designated as NP-hard. The problem is hard because of the necessity to consider exhaustive input permutations in order to realize potential matches, and the number of possible input permutations grows exponentially with the number of inputs. It is also known that the problem becomes much harder if input complementation is allowed in addition to permutation during the mapping process.

Previous work related to the unification problem has primarily addressed the issue of efficient matching of target functionality with the pre-existing library of standard-cells after considering potential permutations and complementation of inputs.

Present techniques to solve this matching problem use functional signatures to eliminate the necessity for the exhaustive generation of permutations and to limit the search space to a manageable size. Functional signatures capture essential properties of a given function (or library cell) that remain invariant even under input permutations/complementation. For example, the size of the ON-set of a function remains invariant under a permutation of input variables. Several other types of functional signatures are also known.

There also exists related work in the area of functional verification that attempts to match portions of a digital circuit (specification) with portions of another digital circuit (implementation) with an objective of comparing their functional equivalences. Techniques used in this area rely on both structural and functional information in order to identify similar portions of the specification and the implementation. In contrast with technology mapping techniques, functional verification techniques can use structural information since the structural views for both the implementation and the specification are known and available.

Characterization refers to the process by which standard cells designed for use in digital ICs designs are extensively analyzed for a number of characteristics such as, but not limited to, timing, power, area, noise, slew, load capacitance, drive strength, and footprint (i.e., size).

Most of the previous work in the area of characterization has focused on methods for reducing simulation times required for characterization. The methods used were typically based on improved models of the underlying devices and better numerical solution methods to solve the mathematical equations formed by modeling the underlying devices.

Prior work on improving the characterization accuracy of the design primarily concentrated on either (1) improving the gate-level timing analysis by considering such issues as false-paths or (2) by performing a characterization of the entire design at the transistor-level for a small set of all possible inputs (since it is not computationally feasible to completely characterize large designs at the transistor-level).

Much of the previous work related to characterization has focused on reducing the time required for extensive circuit simulation. A major disadvantage of previous efforts in the area of characterization is the assumption that the cells are typically designed and characterized without an idea of their potential points of instantiation in any design.

SUMMARY OF THE INVENTION

The present invention discloses methods for effectively solving all of the problems mentioned above for both traditional design methodologies (e.g., custom, semi-custom, gate-array, standard-cell, FPGA), as well as for newly developed and emerging IC design methodologies. The present invention also provides a method for accurately characterizing the overall IC design based on creating and accurately characterizing large regions of the

design at the transistor-level.

The present teachings include a method for improving the utilization of existing cells in a standard-cell library, minimizing the number of new standard-cells created to implement a digital circuit when on-the-fly creation of standard cells is allowed, and improving the characterization of standard-cells given their context-of-use or environment. Also provided is an accurate method for characterizing the entire design based on creating and accurately characterizing groups of standard-cells at the transistor-level.

In accordance with the present teachings, the method and system of minimizing the number of unique standard cells required to implement the digital circuit is referred to herein as unification. As used herein, characterization refers to a process by which a logic cell is analyzed extensively for its characteristic properties, such as pin-to-pin rise/fall times and power consumption. The characterization aspects of the present teachings are applicable over a wide variety of environmental conditions, such as the various semi-conductor process nodes at which the cell can be fabricated.

Conventional standard cell-based design flows involve the creation of a fixed set of standard-cells (called a library), and their use (instantiation) in implementing multiple designs. Recent efforts to improve design performance have resulted in a new design flow in which the generation of cells with entirely new functionality is performed on-the-fly, i.e., in a dynamic manner. In this context, the following considerations are relevant: minimization of the number of unique new standard-cells that need to be generated; the reduction in the resources required for characterizing the new standard-cells; and the accuracy of results achieved by the characterization process.

In contrast to other techniques related to the unification problem, the present teachings disclose a method and system for identifying matching cells based on at least the following set of criteria: target functionality of the required cell, considering potential

permutations and complementation of inputs; and target constraints on the properties of the required cell such as pin-to-pin rise/fall times, input/output capacitances, switching capacitances, power, foot-prints, noise margins, pin-placement, etc. .

5 In contrast to other characterization techniques, the present invention provides effective utilization of restrictions/information regarding the operating environment (i.e., context-of-use) of the standard-cell(s) being characterized. Such restrictions/information may be derived from context-specific knowledge, as is true for the new design flows that generate cells on-the-fly, where the points of instantiation of any cell in a design are known. This
10 knowledge is used to reduce the resources required for characterizing new standard-cells and for improving the accuracy of the characterization process. In addition to the design environment knowledge, constraints utilized by the present invention may come from other data concerning the cell and its intended use.

15 The above and other objects, advantages, and benefits of the present invention will be understood by reference to following detailed description and appended sheets of drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

20 Fig. 1 is a depiction of a standard-cell for implementing a logic function;

Fig. 1a is table listing the functional matches for the standard-cell of Fig. 1;

25 Fig. 2 is a depiction of a logic cell for implementing a logic function, including a timing chart;

Fig. 2a is table listing the matches for the cell of Fig. 2, taking into account both function and timing;

Fig. 3 is a depiction of a logic cell for implementing a logic function, including a timing chart listing timing variances;

5 Fig. 3a is table listing the matches for the cell of Fig. 3, taking into account function, timing, and timing variance;

Fig. 4 is an exemplary depiction of a logic cell for implementing a logic function, including a timing chart listing timing variances;

10

Fig. 4a is a depiction of a logical flow diagram for matching cell function and constraints, including a sorted list of the constraints of Fig. 4;

15

Fig. 5 is a flowchart of an exemplary method for the matching process of the present invention;

Fig. 6 depicts how a context of use can restrict inputs combinations necessary to characterize a combinational cell Z having three inputs, a, b, and c;

20

Fig. 7 further demonstrates the use of context of the cell Z of Fig. 6, to minimize complexity of characterization;

25

Fig. 8 is an exemplary illustration of the greater accuracy of transistor-level characterization, compared to gate-level static timing analysis (STA) in an exemplary circuit as shown; and

Fig. 9 is an exemplary abstraction of characterizing and optimizing clusters of standard cells in accordance with the present invention: the circuit has a critical path through the cells A1-A4-A8-A10-A11.

DETAILED DESCRIPTION OF THE INVENTION

To reiterate, recent efforts to improve design performance have resulted in a new IC design flow in which the generation of cells with entirely new functionality is performed on-the-fly. These cells may be introduced into an existing library of cells, thereby complying with conventional design flow methods. Moreover, such new design flows attempt to better utilize existing cells in a given standard-cell library, even if conventional synthesis processes fail to utilize existing cells — as is commonly the case.

In order to more effectively implement such design flows, procedures for performing context-sensitive and constraint-driven uniquification, and context-sensitive characterization become important. It should be noted that the library of logic cells in the IC design flows considered herein may include cells added during the optimization phase of the original library before the start of the uniquification and characterization procedures of the present invention.

As previously discussed, prior attempts have been made to address the problem of matching a target functional description with the functionality of a physical implementation. Methodologies employed by prior approaches to perform such matching were illustrated in connection with the discussion of Fig. 1. Additionally, there exists related work that attempts to match portions of a digital circuit (specification) with portions of another digital circuit (implementation) with an objective of comparing their functional equivalences, utilizing both structural and functional information.

In contrast with heretofore methodologies in the area of digital circuit specification and implementation matching, the present invention addresses the problem in the context of constraints. Thus, in addition to functionality, the target for matching is annotated with constraints that preferably relate to the target design's use in a design environment.

Since one of the goals of these teachings is to provide support for design flows where cells are generated on-the-fly, the target cell in the library preferably satisfies both the functionality (subject to any permutations/complementations of the cell inputs) and the cell's target constraints. Constraints specified as part of the target may be related to, for example, timing, power, area, noise margins, slew, input/output capacitances, switching capacitances, drive strength, footprint size, pin-placement, etc.

It should be appreciated that the addition of constraints most often decreases the potential of any matches in a specific (i.e., given) library of cells. Referring to Fig. 2 there is shown an example for the matching problem of Fig. 1, with added constraints 220 on the rise and fall timing on the pin-pair z-f. Specifically, the worst z-f rise transition is required to occur in exactly 0.15 ns and the worst z-f fall transition is required to occur in exactly 0.12 ns. These constraints reduce the potential matches from two as shown in Fig. 2a at 205 and 210, if only functional information is considered, to one, i.e., (a:y, b:z, c:x) 215. This is the only match that satisfies the functionality and the timing constraints.

The timing constraints of the example of Fig. 2 were exact. However, during optimization, it is often the case that the target constraints may not be exact. For example, it might be acceptable for a cell to match the target timing constraint within a certain bound or variance. In fact, most constraints are naturally expressed as values within a range. For example, a timing constraint may be expressed as having a rise time that should be no worse than 0.23 ns.

In order to capture variances in timing, the present invention proposes that a constraint is preferably expressed in terms of a nominal value and associated range. For example, consider Fig. 3 that includes the concept of a variance to the example of Fig. 2. Each input pin to the output pin has a rise and a fall time constraint expressed in terms of a nominally required value and a variance as shown in a cell timing chart 310. Cell timing chart 310

includes a lower bound on the variance and an upper bound on the variance. For example, the constraint on pin x-f can now be interpreted as having a rise time delay that is preferred to be 0.20ns, but is no worse than 0.21 (i.e., $0.20 + 0.01$)ns and is no smaller than $-\infty$ ($0.20 - \infty$)ns. With these criteria, the only potential match satisfying the functionality and constraints
 5 in this example is that of (a:y, b:z, c:x), 335.

Similar to the timing constraints and variances illustrated in this example, it is possible to define and use other constraints and variances for other objectives as discussed earlier. Thus, the exemplary discussions presented herein should not be viewed as limitations
 10 of the present teachings.

In a departure from other work related to the area of unification, the present teachings incorporate and use variances on constraints for matching targets with existing cells in the library.
 15

Since the satisfaction of constraints is a key aspect of the present teachings, another aspect hereof is to develop signatures that capture and filter out impossible matches from the existing library. Consider the example of Fig. 4 that illustrates the construction and use of signatures based on constraints for an example target objective of timing 420.
 20

The signature for the target objective of timing can include, but not necessarily or limited to, a sorted set of rise and fall times. It is important to note that such a sorted list is preferably independent of any permutations on input pins. It is also possible to combine rise and fall times and construct signatures comprised of sorted lists of functions of rise and fall
 25 times that would also remain independent of any permutation or potential complementations at the inputs of the library cell (e.g., 405).

In order to account for variances on constraints when computing rise times the signature includes the worst possible rise time (nominal + worst tolerance), and when

computing fall times the signature includes the worst possible fall time (nominal + worst tolerance). Similar lists can be constructed with the sorted rise and fall times for the existing library cells and can be matched against the target list for any potential feasible matches. The feasibility is checked by ensuring that for every target constraint value on a pin (e.g., rise
 5 time), there exists at least one value on an actual pin that satisfies the constraint.

Additional procedures are preferably employed to check that multiple constraints will be matched by the same available, actual, pin data, as depicted in Fig. 4. In steps 430 and 435, the rise time (nominal + worst tolerance) and fall time (nominal + worst tolerance)
 10 constraints are accounted for, respectively, to obtain sorted lists L1 and L2. Sorted lists L1 and L2 are evaluated for feasibility in step 440. Thereafter, the feasibility determination of step 440 is noted (i.e., stored or recorded) in step 445.

The flowchart of Fig. 5 shows an exemplary method for the matching process of the
 15 present invention. Flow chart 500 illustrates how multiple signatures can be constructed based on constraints and used for filtering existing library cells and elements. The signature checking phase is performed in steps 510, 515, 520, and 525. Once the signature checking phase is complete, as evidenced by a “no” determination at step 510, exhaustive permutations and complementations are generated at step 535. A determination is next made whether there
 20 is an existing cell in the library that satisfies the target functionality and constraints at step 540. The permutations and complementations are exhausted by iteratively performing steps 535, 540, 545, and 550. Once it is determined that the signature checking phase will not yield any additional matches, the matching process is concluded at step 560.

25 Note that process 500 allows multiple matches to be generated. In the presence of multiple matches, selection criteria can be utilized to select the most appropriate match. Such criteria may include, for example, preferences for certain constraints and how closely the cell matches the given constraints (based on the nominal values).

Again, characterization refers to the process by which standard cells designed for use in digital IC designs are extensively analyzed for a number of characteristics such as, but not limited to, timing, power, area, noise, slew, load, capacitance, drive strength, and footprint. The present invention provides a method and system of characterization wherein cells are
5 preferably characterized on an as needed basis.

That is, in accordance with the present invention, the characterization process is applied to the range of characteristics relevant to the specific context in which the cell is to be used. As such, the characterization aspect of the present invention provides, for example, the
10 benefits of more accurate characterization, and characterization requiring less (computational) resources.

Fig. 6 illustrates a cell Z (600) having three inputs, a, b, and c. In a conventional design methodology, cell Z can be instantiated in any portion of the design. Hence, any
15 characterization method that does not have knowledge of the context in which a cell is to be used has to consider the possibility that any of the possible input combinations are valid.

In contrast, consider the scenario that occurs in the alternative design flows where standard cells are generated on-the-fly. In this scenario (and other possible scenarios), the
20 design environment for a cell might be known. The design environment is defined by the places of instantiation of the cell in the design. Hence, certain combinations of input values may not be possible for application to cell 600 as illustrated in Fig. 6. As shown, only a subset 620 of possible inputs 605 need be considered due to the design context of cell 600.

Therefore, in accordance with the teachings herein, the characterization process
25 performed on the cell may not need to consider all possible input combinations in order to determine the relevant properties for that cell.

Fig. 7 depicts an exemplary case that illustrates the savings possible by the context-

sensitive characterization of the present invention. A goal is to characterize the rise and fall timing behavior for pin pair c-f 705. The typical characterization process applies all possible (0-1) and (1-0) input transitions at input c for all possible combinations of values at the other inputs (a, b) and obtains the worst-case rise and fall times noticed at the output pin f 710.

5

In accordance with the teachings herein, when the functionality of the cell is known it is also possible to refine the set of characterization vectors. See for example, cel 715 and its associated possible inputs. In fact, the following exemplary coded programming function generates a minimal set of test vectors that provides a complete characterization of a given standard-cell implementing a combinational function:

10

```

GenerateSmartVectors(function)
{
    // supportSize represents the number of variables
    // that function depends on
    supportSize = function.SupportSize();
    for (i = 0; i < supportSize; i++) {

        // compute boolean difference w.r.t input i
        boolDiffFunction = function.BooleanDifference(i);

        setOfVectors = boolDiffFunction.EnumerateMinterms();

        // originalNumVectors += power(2, (supportSize-1));
        // newNumVectors += (boolDiffFunction.CountMinterm())/2;
    }
}

```

15

20

25

Furthermore, if the cell were in the design context as identified in 605 of Fig. 6, it

would suffice to apply only the transitions that are indicated in 715, thus saving valuable time and resources required for the characterization process. In order to achieve such savings, it is important to identify the complete set of properties that are necessary for the operation of the cell in its environment. The cell may be further characterized for a different environment
 5 having different restrictions.

It should be appreciated the methods herein can be extended to encompass (1) restrictions on the characterization process derived from sources other than those discussed above, including, for example, the number of semi-conductor process nodes at which the cell
 10 needs to be characterized; (2) restrictions in any form that improves the accuracy or efficiency of the characterization process, for example, a restriction on the number of temperature settings at which simulation needs to be run, so that the characterization process can be completed faster even if the number of vectors stays the same; and (3) characterization used to determine to all relevant properties of the cell, including for example, the power
 15 characteristics of the cell.

An exemplary circuit 800 is shown in Fig. 8 to illustrate the benefits that can be gained from transistor-level characterization of portions of a digital logic circuit as compared to gate-level static timing analysis (STA). Circuit 800 represents a four-stage chain of NAND
 20 gates. Table 805 depicts the worst propagation delays through the design as determined by both gate-level STA and transistor-level timing analysis for five circuits. Each row in table 805 represents one unique circuit configuration.

The first column represents the drive strengths of the gates at each stage, where zero
 25 (0) represents the smallest drive strength and four (4) represents the largest drive-strength. The second column shows the worst propagation delay through the design as determined by gate-level (STA). The third column represents the worst propagation delay through the same design as determined by transistor-level timing analysis.

Note that the results are based on state-of-the-art process and delay models; and the same load capacitance and input slew values have been used for each of the runs.

It is clear from the data of table 805 that transistor-level timing analysis is much more accurate than gate-level timing analysis. In fact, the differences in timing analyses can be dramatically different in certain kinds of certain configurations as seen in the first circuit configuration (i.e., the first row in the table 805) where each NAND gate is of drive strength 0.

The extent of differences vary based on factors such as the delay models being used, circuit topology, model extraction technology, transistor simulation techniques being used, process in use, circuit design style and drive-strengths of the circuit elements under consideration. Gate-level STA is made inherently conservative to account for potential inaccuracies in the creation of the gate-level abstraction from the transistor-level circuit. Since gate-level analysis is more conservative than transistor-level analysis, it is possible to improve the accuracy of analysis of performance of a design by considering clusters (i.e., partitions or groups) of standard-cells and analyzing and optimizing the clusters at the transistor-level.

A simplified, yet exemplary abstraction of optimizing clusters in accordance with the present invention is shown in Fig. 9. Circuit 900 includes standard cells A1, ..., A12 with a critical path through the cells A1-A4-A8-A10-A11.

Circuit 900 is analyzed at the transistor-level in the areas of interest, i.e., on critical regions, by creating larger cells C1, C2 and C3, as shown in Fig. 9b at 905, 910, and 915, respectively. Cells C1, C2 and C3 are characterized at the transistor-level, and gate-level STA is performed using the new cells. This characterization results in a more accurate (and hence more aggressive!) timing analysis of the design under consideration.

An important aspect of this process is that larger standard-cells, corresponding to C1, C2 and C3 are created, characterized and modeled to comply with cell-level design methodology. It is also important to note that the number of intra-cell wires (i.e., interconnects) is smaller in circuit 920 than in circuit 900. This results in reduction in the number of non-characterized interconnects in the design. In fact, if computational resources were available to perform a detailed transistor-level analysis of the entire circuit 900 under consideration, one could get the most accurate analysis of the timing characteristics of the design. However, considering the practical limitations of computational resources, it is typically feasible to characterize only a few hundred new cells having 20 - 40 transistors per cell, in a reasonable amount of time such as 1-2 days.

In another aspect of the present teachings, it is noted is that the characterization method herein is beneficial even when the larger cells can be optimized by creating entirely new transistor structures that are functionally equivalent to the standard-cells that they replace.

It should also be appreciated by those skilled in the art that the particular method and systems of context-sensitive constraint driven unification and characterization and other aspects of the teachings herein are but examples of the present invention. For example, the methods disclosed herein may be implemented in a system including a computer programmed to execute program instructions for effectuating the IC design methods of the present invention. The program instructions may be embodied on a storage medium such as a hard disk drive, a CD-ROM, or any other computer readable media. Thus, the particular method and systems described herein do not limit the scope or variety of applications that the present invention may be suitably implemented. Thus, it should be understood that the foregoing description is only illustrative of a present implementation of the teachings herein. Various alternatives and modification may be devised by those skilled in the art without departing from the invention. Accordingly, the present invention is intended to embrace all such alternatives, modifications, and variances that fall within the scope of the appended claims.